

REMARKS

This responds to the Office Action dated on June 11, 2007.

Claim 7 is amended, no claims are canceled, and no claims are added; as a result, claims 1-27 are now pending in this application.

§102 Rejection of the Claims

Claims 1-4, 7, 8, 10-15, 18-19, 21-23, 26 and 27 were rejected under 35 U.S.C. § 102(e) as being anticipated by Forin (US Pat: 6,321,276).

Applicant's invention as claimed in claims 1, 12, 22 and 23 pertains to a method and apparatus for node translation in a clustered multiprocessor system. The method and apparatus supports user-level communications without the need for OS(operating system) intervention on communication events between a source (master) node and a destination (slave) node. Specifically, Applicant teaches, and claims in claims 1, 12, 22 and 23, using a communication engine including a local connection table (LCT) configured to locate the source and the destination nodes by using corresponding connection descriptors (CDs). Forin does not teach or suggest the local connection table (LCT) using the connection descriptors (CDs) as taught by Applicant and claimed in claims 1, 12, 22 and 23.

By locating the participating nodes through the LCT using the CDs, the node translation method under Applicant's approach allows user programs to use the same access mechanism for intra-machine and inter-machine communications. In addition, since the node translation between a local node and a remote node is defined in user programs without the need for OS intervention as discussed above, the node translation method under Applicant's approach reduces the overhead in communication between a local node and a remote node. For example, as discussed in the specification at p. 2, lines 16-28 and p. 3, lines 17-26, under the traditional approach, if the virtual address range spans 15 physical pages in a remote node, the OS first checks whether it has address translations for all of the virtual addresses, and then generates 15 separate queued block-transfer requests to the remote node to cause 15 separate physical transfers. Instead, under Applicant's approach, user programs can directly and simultaneously define the multiple physical page transfers to a remote node using the LCT and the CDs and

thereby reduce the overhead caused by the OS intervention. See p. 2, lines 16-28 and p. 3, lines 17-26 of the Specification.

To accomplish this, Applicant teaches and claims, for example, in claim 1, “assigning a connection descriptor to a virtual connection, the connection descriptor being a handle that specifies an endpoint node for the virtual connection,” “defining a local connection table accessible by the communication engine, wherein the local connection table is configured to be accessed using the connection descriptor to produce a system node identifier for the endpoint node,” “accessing, via the communication engine, the local connection table using the connection descriptor of the communication request to produce the system node identifier for the endpoint node for the virtual connection,” and “sending a memory request from the communication engine to the endpoint node, wherein the memory request is sent to the local processing element node if the endpoint node is the local processing element node, and is sent over the network interconnect to the remote processing element node if the endpoint node is the remote processing element node and wherein transfer of data associated with the memory request occurs independently of the processor that generated the communication request.” Claims 12, 21 and 23 also have corresponding limitations.

Forin describes a method and system for processing input/output requests between a host computer and its locally associated I/O devices to recover from a local translation failure. When a translation failure occurs in the ‘recoverable I/O request processor’ in an I/O device (Fig. 3, block 60 and Fig. 6, block 60a), the ‘recoverable I/O request processor’ requests virtual address mapping information via the ‘virtual memory manager’ in a host computer (Fig. 3, block 64) from the operating system associated with the host computer.

Forin does not, however, teach or suggest each and every element in claim 1. First of all, col. 3, lines 40-65 and col. 6, lines 47 through col. 7, lines 19 in Forin describe a descriptor for communicating an I/O request from a host computer to an I/O device locally associated with the host computer. See also Figs. 1-3, 6 and 6(a). The descriptor may include “a control field for storing a code indicative of an I/O operation,” “a buffer virtual address field for storing a virtual memory address of a buffer to be utilized in an I/O operation,” “a translation error control flag field for storing a translation error control flag for instructing a recoverable I/O request processor of an I/O device to attempt to recover from a virtual address translation failure.” There is no

indication, however, that the descriptor in Forin is used to specify the identification of a remote node involved in an I/O device request. Also, contrary to the Examiner's assertion, the "logical connection" in Forin merely describes a communication network (e.g., a LAN or a WAN). We cannot assume that Forin's descriptor for an I/O request includes an identifier for a remote host computer. The cited portions, therefore, do not teach or suggest "assigning a connection descriptor to a virtual connection, the connection descriptor being a handle that specifies an endpoint node for the virtual connection" as claimed in claim 1.

Col. 14, lines 52-67 in Forin describes recovering from virtual address translation failures using the Virtual Interface Architecture (VIA) network interface adapter (Fig. 6, block 100) including the recoverable I/O request processor (block 60a). The recoverable I/O request processor translates virtual memory addresses to physical memory addresses, maintains local page tables, and recovers from local address translation failures by requiring virtual mapping information via the virtual memory manager (Fig. 3, block 64) from its associated OS. The cited portion shows that the local page tables in Forin are mapping information between virtual and physical addresses for the physical pages associated with the local host computer. In contrast, as taught by Applicant at Figs. 4A and 4B, the local connection table (LCT) (block 104) under Applicant's approach is a table separate from a local page table (SHUB TLB, block 108). The LCT contains identifications not only for a local node but also for a remote node. Applicant is unable to find these teaching in Forin. The cited portion, therefore, does not teach or suggest "defining a local connection table (LCT) configured to be accessed using the connection descriptor (CD) to produce a system node identifier for the endpoint node" as taught by Applicant and claimed in claim 1.

Although Col. 15, lines 19-36 in Forin states that "the virtual interface (VI) kernel agent (Fig. 6, block 105) may establish and break connection with remote machines," the cited portion is not clear as to how the VI kernel agent accomplishes the connection with the remote machines. In contrast, Applicant explicitly teaches and claims in, for example, claim 1, "accessing, via the communication engine, the local connection table (LCT) using the connection descriptor of the communication request to produce the system node identifier for the endpoint node." Applicant is unable to find such a teaching in Forin.

Finally, the Examiner also states that Forin teaches “memory request being sent over the network interconnect to the remote processing element node if the endpoint node is the remote processing element node.” As support of this, the Examiner cites col. 14, lines 32-51 in Forin, the part of which states:

(col. 5, line 41-51) ... In order to improve I/O performance, VIA (Virtual Interface Architecture) network interface adapters are capable of performing local virtual memory address translations. However, conventional VIA network interface adapters are incapable of recovering from local translation failures. The behavior of a conventional VIA network interface adapter typically is a fatal error if such a condition occurs, which may include breaking connection with a remote host. The present embodiment enhances the I/O performance of VIA network interface adapters by **enabling the VIA network interface adapter to recover from local translation failures.**

As quoted above, the cited portion describes that, for example, “if the recoverable I/O request processor (60) is unable to translate a virtual memory address because address translation tables local to the device do not contain a valid entry for the virtual memory address of interest, the recoverable I/O request processor is preferably capable of obtaining the virtual memory mapping information from an external source, e.g., from the operating system associated with its locally connected host computer via the virtual memory manager (64), without failing the I/O operation. See also Abstract, Figs. 2-3 and col. 8, lines 19-29. There is, however, no sending the I/O request to a remote node when it is indicated that the virtual memory is associated with a remote node. Instead, when a local translation fails (e.g., the translation address is for a virtual address in a remote node and fails in the local translation table), the recoverable I/O request processor asks for the corresponding mapping information from its locally associated operating system. This, therefore, teaches away from the sending the memory request to the remote node as taught by Applicant and claimed in claim 1.

For these reasons, Forin does not teach or suggest the method and apparatus for node translation method in a clustered multiprocessor system using a communication engine including a local connection table (LCT) configured to locate the source and the destination nodes by using corresponding connection descriptors (CDs) as taught by Applicant and claimed in claims 1. Similar arguments can be applied to other independent claims 12, 22 and 23. Reconsideration is respectfully requested.

With regard to claims 2-4, 7, 8, 10, 11, 13-15, 18, 19, 21, 26 and 27, claims 2-4, 7, 8, 10, 11, 13-15, 18, 19, 21, 26 and 27 are patentable as dependent on a patentable base claim.

Allowable Subject Matter

Claims 5, 6, 9, 16, 17, 20, 24 and 25 are objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent from including all of the limitations of the base claim and any intervening claims.

Applicant respectfully submits that the base claims are patentable but reserves the option to recast these claims as independent claims, if necessary.

RESERVATION OF RIGHTS

In the interest of clarity and brevity, Applicant may not have addressed every assertion made in the Office Action. Applicant's silence regarding any such assertion does not constitute any admission or acquiescence. Applicant reserves all rights not exercised in connection with this response, such as the right to challenge or rebut any tacit or explicit characterization of any reference or of any of the present claims, the right to challenge or rebut any asserted factual or legal basis of any of the rejections, the right to swear behind any cited reference such as provided under 37 C.F.R. § 1.131 or otherwise, or the right to assert co-ownership of any cited reference. Applicant does not admit that any of the cited references or any other references of record are relevant to the present claims, or that they constitute prior art. To the extent that any rejection or assertion is based upon the Examiner's personal knowledge, rather than any objective evidence of record as manifested by a cited prior art reference, Applicant timely objects to such reliance on Official Notice, and reserves all rights to request that the Examiner provide a reference or affidavit in support of such assertion, as required by MPEP § 2144.03. Applicant reserves all rights to pursue any cancelled claims in a subsequent patent application claiming the benefit of priority of the present patent application, and to request rejoinder of any withdrawn claim, as required by MPEP § 821.04.

CONCLUSION

Applicant respectfully submits that the claims are in condition for allowance, and notification to that effect is earnestly requested. The Examiner is invited to telephone Applicant's attorney at (612) 373-6909 to facilitate prosecution of this application.

If necessary, please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

Respectfully submitted,

STEVEN L. SCOTT (DO NOT USE) ET AL.

By their Representatives,

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

P.O. Box 2938

Minneapolis, MN 55402

(612) 373-6909

Date Sept. 11, 2007

By Thomas F. Brennan
Thomas F. Brennan
Reg. No. 35,075

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: Mail Stop Amendment, Commissioner of Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on this 11 day of September 2007.

Lisa Rosarske

Name

Sven Rosarske

Signature